To See or Not to See: The Need for Attention to Perceive Changes in Scenes

Ronald A. Rensink, J. Kevin O'Regan¶ & James J. Clark‡

Cambridge Basic Research

Nissan Research & Development, Inc.

Cambridge, Massachusetts, USA

¶*Laboratoire de Psychologie Experimentale, CNRS*
 Université René Descartes,
 Paris, France

‡*Department of Electrical Engineering,*
 McGill University,
 Montreal, Quebec, Canada

Running Head: The Need for Attention to See Change

Author Notes

Abstract

When looking at a scene, our impression as observers is that we see its entire structure in great detail and can immediately notice any changes in it.   However, when brief blank fields are placed between alternating displays of an original and a modified scene, a striking failure of perception is induced: identification of changes becomes extremely difficult, even when changes are large and made repeatedly.   Identification is much faster when a verbal cue is provided, showing that poor visibility is not the cause of this difficulty. Identification is also faster for objects mentioned in brief verbal descriptions of the scene. These results support the idea that observers never form a complete, detailed representation of their surroundings. In addition, results also indicate that attention is required to perceive change, and that in the absence of localized motion signals it is guided on the basis of high-level "interest".

To See or Not to See: The Need for Attention
to Perceive Changes in Scenes

Although we must look in order to see, looking by itself is not enough. For example, if we turn our eyes toward a bird singing in a tree, we will often fail to see it right away, "latching onto" it only after some effort. This also holds true for objects in plain view: if our mind wanders during driving, we can often miss important road signs, even when these are highly visible. In both situations, the information needed for perception is available to the observer. Something, however, prevents observers from using this information to see the new objects that have entered their field of view.

In this paper it is argued that the key factor is attention. In particular, it is proposed that the visual perception of change in a scene occurs only when focused attention is given to the part being changed. In support of this position, it will be shown that when the low-level cues that draw attention are swamped, large changes in images of real-world scenes become extremely difficult to identify, even though these changes are repeated dozens of times and observers have been told to expect them. Changes are easily identified when a valid verbal cue is given, indicating that stimulus visibility is not reduced. Changes are also easily identified when made to objects mentioned in brief verbal descriptions of the scene. Taken together, these results indicate that—even when sufficient viewing time has been given—we do not build up a representation of a scene that allows us to automatically perceive change. Rather, perception of change is mediated through a narrow attentional bottleneck, with attention attracted to various parts of a scene based on high-level "interest".

The phenomenon of "induced change blindness" has previously been encountered in two rather different experimental paradigms. The first, concerned with visual memory, investigated the detection of change in briefly-presented arrays of simple figures or letters (e.g., Pashler, 1988; Phillips, 1974). Here, an initial display was presented for 100-500 ms, followed by a brief interstimulus interval (ISI), followed by a second display in which one of the items was removed or replaced on half the trials. Responses were forced-choice guesses about whether a change had occurred. Observers were found to be poor at detecting change if old and new displays were separated by an ISI of more than 60-70 ms.

The second type of paradigm, stemming from eye-movement studies, examined the ability of observers to detect changes in an image made during a saccade (e.g., Bridgeman,

Hendry & Stark, 1975; Grimes, 1996; McConkie & Zola, 1979). A variety of stimuli were tested, ranging from arrays of letters to images of real-world scenes. In all cases, observers were found to be quite poor at detecting change, with detection good only for a change in the saccade target (Currie, McConkie, Carlson-Radvansky & Irwin, 1995).

Although blindness to saccade-contingent change has been attributed to saccade-specific mechanisms, the blurring of the retinal image during the saccade also masks the transient motion signals that normally accompany an image change. Since transients play a large role in the drawing of attention (e.g., Klein, Kingstone & Pontefract, 1992; Posner, 1980), saccade-contingent change blindness may not be due to saccade-specific mechanisms, but rather may originate from a failure to correctly allocate attention. The blindness to changes in briefly-presented displays may have a similar cause: in those experiments, detection was not completely at chance, but corresponded to a monitoring of 4-5 randomly-selected items, a value similar to the number of items that can be simultaneously attended (Pashler, 1987; Pylyshyn & Storm, 1988; Wolfe, Cave & Franzel, 1989).

In order to examine whether both types of change blindness might be due to the same attentional mechanism, and whether this mechanism might also lead to change blindness under more normal viewing conditions, a "flicker" paradigm was developed. Here, an original image A repeatedly alternates with a modified image A', with brief blank fields placed between successive images (Fig. 1). Differences between original and modified images can be of any size and type; in the experiments presented here the changes are chosen to be highly visible. The observer freely views the flickering display and hits a key when the change is perceived. In order to prevent guessing, the observer must then correctly report the type of change and describe the part of the scene that was changing.

This paradigm allows the ISI manipulations of the brief-display techniques to be combined with the free-viewing conditions and perceptual criteria of the saccade-contingent methods. And because the stimuli are available for long stretches of time and no eye movements are required, it also provides the best opportunity possible for an observer to build a representation conducive to perceiving changes in a scene. The change blindness found with the brief-display techniques might have been caused by insufficient time to build an adequate representation of the scene; saccade-contingent change blindness might have been caused by disruptions due to eye movements. Both of these factors are eliminated in the flicker paradigm, so that if they are indeed the cause of the difficulties, perception of

change should now become easy. But if attention is the key factor, a different outcome will be expected. The flicker caused by the blank fields will swamp the local motion signals due to the image change, preventing attention from being drawn to its location. Observers will then fail to see large changes under conditions of extended free viewing, even when these changes are not synchronized to saccades.

General Method

In the experiments reported here, flicker sequences were usually composed of an original image A and modified version A' displayed in the sequence A, A, A', A',…with gray blank fields placed between successive images (Fig. 1). Each image was displayed for 240 ms and each blank for 80 ms. Note that each image was presented twice before being switched. This created a degree of temporal uncertainty as to when the change was being made, and also allowed for a wider range of experimental manipulation.

Figure 1 about here

All experiments used the same set of 48 color images of real-world scenes. Images were 27º wide and 18º high. A single change—color, location, or presence/absence—was made to an object or area in each. To test for the influence of higher-level factors, changes were further divided according to the degree of interest in the part of the scene being changed. Interest was determined via an independent experiment in which five naive observers provided a brief verbal description of the scene: central interests (CIs) were defined as objects or areas mentioned by three or more observers; marginal interests (MIs) were objects or areas mentioned by none. The changes to the images were such that the average change in intensity and color of the MIs and the CIs were similar, while the areas of the MI changes (average = 22 sq. deg) were somewhat larger than the CI changes (average = 18 sq. deg). In all cases, changes were quite large and easy to see once noticed. For example, a prominent object could appear and disappear, switch its color between blue and red, or shift its position by a few degrees (Fig. 2).

Ten naive observers participated in each experiment. They were instructed to press a key when they saw the change, and then to verbally describe it. Before each experiment, observers were told of the types of change possible, and were given six practice trials (two

examples of each type) to familiarize themselves with the task. Images were presented in random order for each subject. The dependent variable was the average number of alternations [proportional to the reaction time] needed to see the change. Averages were taken only from correct responses, i.e., responses where the observer correctly identified both the type of change occurring and the object or area being changed. As might be expected from the use of large changes, identification error rates were low, averaging only 1.2% across all experiments.

Figure 2 about here

Experiment 1

Experiment 1 examined whether the basic "flicker paradigm" could indeed induce change blindness. Images were displayed for 240 ms and blanks for 80 ms, with images repeated before being switched (Fig. 1). If insufficient viewing time were the reason for the change blindness found in the brief-display experiments, changes should now be seen within at most a few seconds of viewing. If saccade-specific mechanisms were responsible for the change blindness found in the saccade experiments, changes should now be easy to see simply by keeping the eyes still. But if the failures to detect change were due to an attentional mechanism, changes under these flicker conditions should take a long time to see.

The results of Experiment 1 (Fig. 3a) show a striking effect: under flicker conditions, changes in MIs were extremely difficult to see, requiring an average of 17.1 alternations [10.9 seconds] before being identified; indeed, for some images observers required over 80 alternations [50 seconds] to identify a change that was obvious once noticed. Changes in CIs were noticed much more quickly, with an average of 7.3 alternations [4.7 seconds]. Because discriminability was not equated for the three different types of change, performance between them cannot be compared. However, within each type of change, perception of MI changes took significantly longer than CI changes (p < .001 for appear/disappear; p < .05 for color; p < .001 for location), even though MI changes were on average over 20% larger in area.

To confirm that the changes in the pictures were indeed easy to see when flicker was absent, the experiment was repeated with the blanks in the displays removed. A completely

different pattern of results now emerged: identification required only 1.4 alternations [0.9 seconds] on average, showing that observers quickly noticed the changes. No significant differences were found between MIs and CIs for any type of change, and no significant differences were found between types of change (p > 0.3 for all comparisons).

Figure 3 about here

Experiment 2

One explanation for the poor performance found in Experiment 1 might be that old and new scene descriptions could not be compared because of time limitations. Although the blanks between images were only 80 ms—well within the 300 ms duration of iconic memory (e.g., Irwin, 1991; Sperling, 1960)—it has been shown that approximately 400 ms are needed to process and consolidate an image in memory (Potter, 1976). Since the images in Experiment 1 were displayed for only 240 ms, this may have interfered with consolidation, and thus with the ability to compare successive images.

In Experiment 2, therefore, the blanks between pairs of identical images were "filled in" by replacing them with an 80 ms presentation of the "surrounding" images. Thus, instead of presenting each image for 240 ms, followed by a blank for 80 ms, and then presenting it again for another 240 ms, images were now presented without interruption for 240 + 80 + 240 = 560 ms at a time. Because the blanks between the original and modified images were kept, original image A and modified image A' were now presented in the sequence A, A', A, A'…, with changes continuing to be made at the same rate as before. If memory processing were the limiting factor, the longer display of the images should now allow consolidation to take place, and so cause the changes to be much more easily seen.

Results (Fig. 3b), however, show that this did not occur. Although there was a slight speedup for MI changes, this was not large; indeed, response times for MIs and CIs for all three kinds of change were not significantly different from their counterparts in Experiment 1. Note that this also shows that the temporal uncertainty caused by the repeating images in Experiment 1 does not greatly affect performance: pairs of identical images each of duration 240 ms and separated by 80 ms have much the same effect as a single image presented for 560 ms.

Experiment 3

Another possible explanation for the occurrence of change blindness under flicker conditions is that the flicker reduces the visibility of the items in the image to the point where they simply become difficult to see. To examine this possibility, Experiment 3 repeated Experiment 1, but with verbal cues (single words or word pairs) placed in white rectangles for 3 seconds at the beginning of each trial. Two different cueing conditions were used. In the partially-valid condition, cues were divided equally into valid cues (naming the part of the scene changed) and invalid cues (naming some other part). In the completely-valid condition, cues were always valid. If visibility is indeed the limiting factor, no large effect of cueing should occur—the target will simply remain difficult to find. Otherwise, performance should be greatly sped up by valid cues, and be relatively unaffected (or even slowed down) by invalid ones.

As Fig. 3c shows, valid cues always caused identification of both MI and CI changes to be greatly sped up. This speedup was significant both for partially-valid cues ($p < .001$ for MI; $p < .03$ for CI) and completely-valid cues ($p < .001$ for both MI and CI). Indeed, for completely-valid cues, the difference in response times for MIs and CIs declined to the point where it was no longer significant. Note that this latter result indicates that the faster performance for CIs in Experiment 1 is unlikely to be due to the simple "salience" of their features—such a near-equality of search times would hardly be expected if the CIs contained features salient enough to preferentially catch the attention of observers.

In contrast to the speedup for valid cues, invalid cues caused a slight slowdown (although this was not found to be significant). Taken together, then, these results show that observers could readily locate a cued target under flicker conditions, thereby demonstrating that visibility was not a limiting factor.

General Discussion

The preceding experiments show that under flicker conditions observers can take a surprisingly long time to perceive large changes in images of real-world scenes. This difficulty is due neither to a disruption of the information received nor to a disruption of its storage. It does, however, depend greatly on the significance of the part of the scene being changed, with identification fastest for those structures of greatest interest.

We therefore make the following proposal: <u>visual perception of change in an object occurs only when that object is given focused attention; in the absence of such attention, the contents of visual memory are simply overwritten (i.e., replaced) by subsequent stimuli, and so cannot be used to make comparisons</u>. Although it is not yet possible to specify the detailed operation of the attentional mechanisms involved, it is likely that the allocation of attention causes the relevant structures to form "object files" (Kahneman, Treisman & Gibbs, 1992), or at least lets them be entered into a more durable store such as visual short term memory (e.g., Coltheart, 1980; Irwin, 1991; Sperling, 1960) so that comparisons can be made.

In this view, all the effects encountered here can be traced back to the allocation of attention, which is either "pulled" by transient motions or "pushed" by volitional control (e.g., Klein, Kingstone & Pontefract, 1992; Posner, 1980). Under normal conditions, the motion signals resulting from a change draw attention to its location and so allow us to perceive it. When these signals are delocalized by flicker, their influence is effectively removed; attention is then directed entirely by static low-level properties such as feature gradients (Nothdurft, 1992) and high-level volition. If there are no distinct low- or high-level cues (true of most stimuli used here) detection of change will then require a slow, item-by-item scan of the entire image, and so give rise to long identification times. The faster identification of CI changes—in spite of their smaller area—would result from the attraction of attention via the high-level interest of the objects that are changed.

If this view is correct, it points towards tighter connections between lines of research in four rather different areas of vision: eye movements, visual attention, visual memory, and scene perception. For example, the failure to find representations capable of providing automatic detection of change supports the view of eye-movement researchers (e.g., Bridgeman, Hendry & Stark, 1975; Irwin, 1991; McConkie & Zola, 1979) that there simply is no spatiotopic buffer where successive fixations are added, compared, or otherwise combined. Note that although the experiments here did not explicitly address the issue of image addition (superposition), the difficulty in detecting positional shifts suggests that such superposition is unlikely—otherwise observers could simply have looked for instances of "doubled" structures, i.e., where the original and the shifted object were both present side-by-side. In any event, it would appear that much—if not all—of the blindness to saccade-contingent change is simply due to the disruption of the retinal image during a saccade, which causes swamping of the local motion signals that would normally draw attention. A

similar explanation can also account for the change blindness encountered in the brief-display studies, suggesting that a common framework may encompass both of these effects.

The results presented here are also related to findings (Mack, Tang, Tuma, Kahn, & Rock, 1992; Neisser & Becklen, 1975; Rock, Linnet, Grant, & Mack, 1992) that attention is required to explicitly perceive a stimulus in the visual field. In those studies, it was found that observers giving their complete attention to particular objects or events in a scene became "blind" to other, irrelevant objects. This required that observers not suspect that the irrelevant objects would be tested, for even a small amount of (distributed) attention would cause these objects to be perceived. The results here are more robust, in that blindness occurred even when observers knew that changes would be made, and so could distribute their attention over the entire picture array if it would help. Thus, although distributed attention object apparently facilitates the perception of object presence, it does not facilitate the perception of change. Presumably, distributed attention is not sufficient to move a structure from visual memory into the more durable store that would allow the perception of change to take place.

In addition to asserting that attended items are entered into a relatively stable store, our proposal also asserts that unattended items are overwritten by new stimuli that subsequently appear in their location. This latter point is based on the finding that change blindness occurs even when images are separated by an ISI of only 80 ms, a time well within the 300 ms limit of iconic memory; if no such replacement took place, observers could simply have used the superposed images of original and shifted objects to find positional shifts. Such a replacement of unattended items has been proposed to explain metacontrast masking (Enns & DiLollo, 1996), and it is possible that the same mechanism is involved here. In any event, this mechanism implies that two rather different fates await items in visual memory: attended items are loaded into a durable store and are perceived to undergo transformation whenever the stimulus is changed, while unattended items are simply replaced by the arrival of new items, with no awareness that a replacement has occurred.

Finally, the work presented here also suggests a tighter connection between attention and scene perception. To see how this comes about, recall that the attentional scanning for the invalidly-cued targets in Experiment 3 did not speed up identification. This indicates that it is not sufficient to have simply paid attention to the object at some point in the past—the perception of change can occur only during the time that the object is being attended (or at

least during the time the object is held in the limited-capacity durable store). After attention is removed, the perception of change vanishes and the objects again become susceptible to replacement. A similar "evaporation" of attentional effect has also been found to occur in visual search, where items obtain no benefit from having been previously attended (Wolfe, 1996). Thus, just as the perception of a scene is mediated by a rapidly-shifting fovea of limited area, so is it also mediated by a rapidly-shifting attentional mechanism limited in the number of items it can handle at any time.

The limited capacity of this mechanism requires that it be used effectively if a scene is to be quickly perceived. But how can appropriate guidance be given if the scene has not yet been recognized? Previous work has shown that the gist of a scene can be determined within 100-150 ms (Biederman, Mezzanotte & Rabinowitz, 1982; Intraub, 1981; Potter, 1976); it may well be that the gist includes a description of the most interesting aspects, which are then used to guide attention. By measuring the relative speed of perceiving changes to various parts of a scene, it might be possible to determine the order in which attention visits the constituent objects and regions. The resultant "attentional scan path" may prove to be an interesting new tool in the study of scene perception, providing a useful complement to techniques that study eye movements and memory for objects as a function of how well they fit the gist of a scene (e.g., Friedman, 1979; Loftus & Mackworth, 1978). Furthermore, the correlation found here between reaction time and degree of "interest" (as derived from written descriptions) opens up another interesting possibility, namely, that the flicker paradigm can be adapted to determine what non-verbal observers (e.g., animals and young children) find interesting in the world.

Why can we look at but not always see objects that come into our field of view? The evidence presented here indicates that the key factor is attention, without which observers are blind to change. Since attention can be concurrently allocated to only a few items (e.g., Pashler, 1987; Pylyshyn & Storm, 1988; Wolfe, Cave & Franzel, 1989), this implies that only a few changes can be perceived at any time. Although such a low-capacity mechanism might seem to be rather limiting, this need not be the case: if it can switch quickly enough so that objects and events can be analyzed whenever needed, little is gained by the simultaneous representation of all their details (Ballard, Hayhoe & Whitehead, 1992; Dennett, 1991; O'Regan, 1992). Thus, given that attention is normally drawn to any change in a scene and is also attracted to those parts most relevant for the task at hand, our impression as observers

will generally be of a richly-detailed environment, with accurate representation of those aspects of greatest importance to us. It is only when low-level transients are masked or are disregarded due to inappropriate high-level control that these representations will not be correctly updated, and so cause the relatively sparse nature of our visual representations to become apparent.

References

Ballard, D.H., Hayhoe, M.M., & Whitehead, S.D. (1992). Hand-eye coordination during sequential tasks. *Philosophical Transactions of the Royal Society of London B,* **337**, 331-339.

Biederman, I., Mezzanotte, R.J., & Rabinowitz, J.C. (1982). Scene perception: Detecting and judging objects undergoing relational violations. *Cognitive Psychology*, **14**, 143-177.

Bridgeman, B., Hendry, D., & Stark, L. (1975). Failure to detect displacements of the visual world during saccadic eye movements. *Vision Research*, **15**, 719-722.

Coltheart, M. (1980). Iconic memory and visible persistence. *Perception & Psychophysics*, **27**, 183-228.

Currie, C., McConkie, G.W., Carlson-Radvansky, L.A., & Irwin, D.E. (1995). Maintaining visual stability across saccades: Role of the saccade target object. *Technical Report UIUC-BI-HPP-95-01,* Beckman Institute, University of Illinois.

Dennett, D.C. (1991). *Consciousness Explained*. Boston : Little, Brown and Co.

Enns, J.T., & DiLollo, V. (1996). Space, time, and the object substitution effect. *International Journal of Psychology*, **31** (3/4), 36.

Friedman, A. (1979). Framing pictures: The role of knowledge in automatized encoding and memory for gist. *Journal of Experimental Psychology: General*, **108**, 316-355.

Grimes, J. (1996). On the failure to detect changes in scenes across saccades. In K. Akins (ed.), *Perception (Vancouver Studies in Cognitive Science, vol. 5).* New York : Oxford University Press.

Intraub, H. (1981). Rapid conceptual identification of sequentially presented pictures. *Journal of Experimental Psychology: Human Perception and Performance*, **7**, 604-610.

Irwin, D.E. (1991). Information integration across saccadic eye movements. *Cognitive Psychology*, **23**, 420-456.

Kahneman, D., Treisman, A., & Gibbs, B. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, **24**, 175-219.

Klein, R., Kingstone, A., & Pontefract, A. (1992). Orienting of Visual Attention. In K. Rayner (ed.), *Eye movements and visual cognition: Scene perception and reading.* New York: Springer.

Loftus, G.R., & Mackworth, N.H. (1978). Cognitive determinants of fixation location during picture viewing. *Journal of Experimental Psychology: Human Perception and Performance*, **4**, 547-552.

Mack, A., Tang, B., Tuma, R., Kahn, S., & Rock, I. (1992). Perceptual organization and attention. *Cognitive Psychology*, **24**, 475-501.

McConkie, G.W., & Zola, D. (1979). Is visual information integrated across successive fixations in reading? *Perception & Psychophysics*, **25**, 221-224.

Neisser, U., & Becklen, R. (1975). Selective looking: Attending to visually significant events. *Cognitive Psychology*, **7**, 480-494.

Nothdurft, H.C. (1992). Feature analysis and the role of similarity in pre-attentive vision. *Perception & Psychophysics*, **52**, 355-375.

O'Regan, J.K. (1992). Solving the "real" mysteries of visual perception: The world as an outside memory. *Canadian Journal of Psychology*, **46**, 461-488.

Phillips, W.A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception & Psychophysics*, **16**, 283-290.

Pashler, H. (1987). Detecting conjunctions of color and form: Reassessing the serial search hypothesis. *Perception & Psychophysics*, **41**, 191-201.

Pashler, H. (1988). Familiarity and visual change detection. *Perception & Psychophysics*, **44**, 369-378.

Posner, M.I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology,* **32**, 3-25.

Potter, M.C. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory*, **2**, 509-522.

Pylyshyn, Z.W., & Storm, R.W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision,* **3***,* 179-197.

Rock, I., Linnet, C., Grant, P., & Mack, A. (1992). Perception without attention: Results of a
new method. *Cognitive Psychology*, **24**, 502-534.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological
Monographs*, **74**, 1-29.

Wolfe, J.M., Cave, K.R., & Franzel, S.L. (1989). Guided search: An alternative to the feature
integration model for visual search. *Journal of Experimental Psychology: Human Perception
and Performance,* **15**, 419-433.

Wolfe, J.M. (1996). Post-attentive vision. *Investigative Ophthalmology & Visual Science*, **37**,
S214.

Figure Captions

FIG. 1.   General design of the flicker paradigm. Trials began with a 3-second gray field containing a white rectangle (in which a word appears when a cue is used). This was followed by a 1-second gray field, followed by a "flicker" sequence that continued until the observer responded or 60 seconds had elapsed. In the example here, original image A (statue with background wall) and modified image A' (statue with wall removed) are displayed in the order A, A, A', A',…with gray fields between successive images.

FIG. 2 Examples of changes in scenes. Here, original and modified images alternated every 640 ms. (a) Change in a marginal interest (MI): railing behind people is moved. Although the railing is easily seen, and its location shift is large (3^o), an average of 16.2 alternations [10.4 seconds] is required for identification. (b) Change in a central interest (CI): helicopter in background is moved. Although the change in location is roughly the same as in the previous case, and the size and contrast of the item changed is comparable, identification requires on average only 4.0 alternations [2.6 seconds].

FIG. 3.   Identification of change under flicker conditions. Error bars indicate one standard error; horizontal gray bars standard error of comparison conditions. (a) Results under the "basic" conditions of Experiment 1. For all types of change, a similar pattern was found: MI changes were difficult to see, whereas CI changes were considerably easier. Dashed line indicates baseline performance when no blanks are present. (b) Effect of longer image duration. Dashed lines indicate results of Experiment 1. As is evident, only a slight speedup occurred. (c) Effects of verbal cues. Valid cues (in both conditions of 50% and 100% validity) caused a large speedup for both MI changes and CI changes. In contrast, invalid cues (in conditions of 50% validity) caused slight slowdowns.
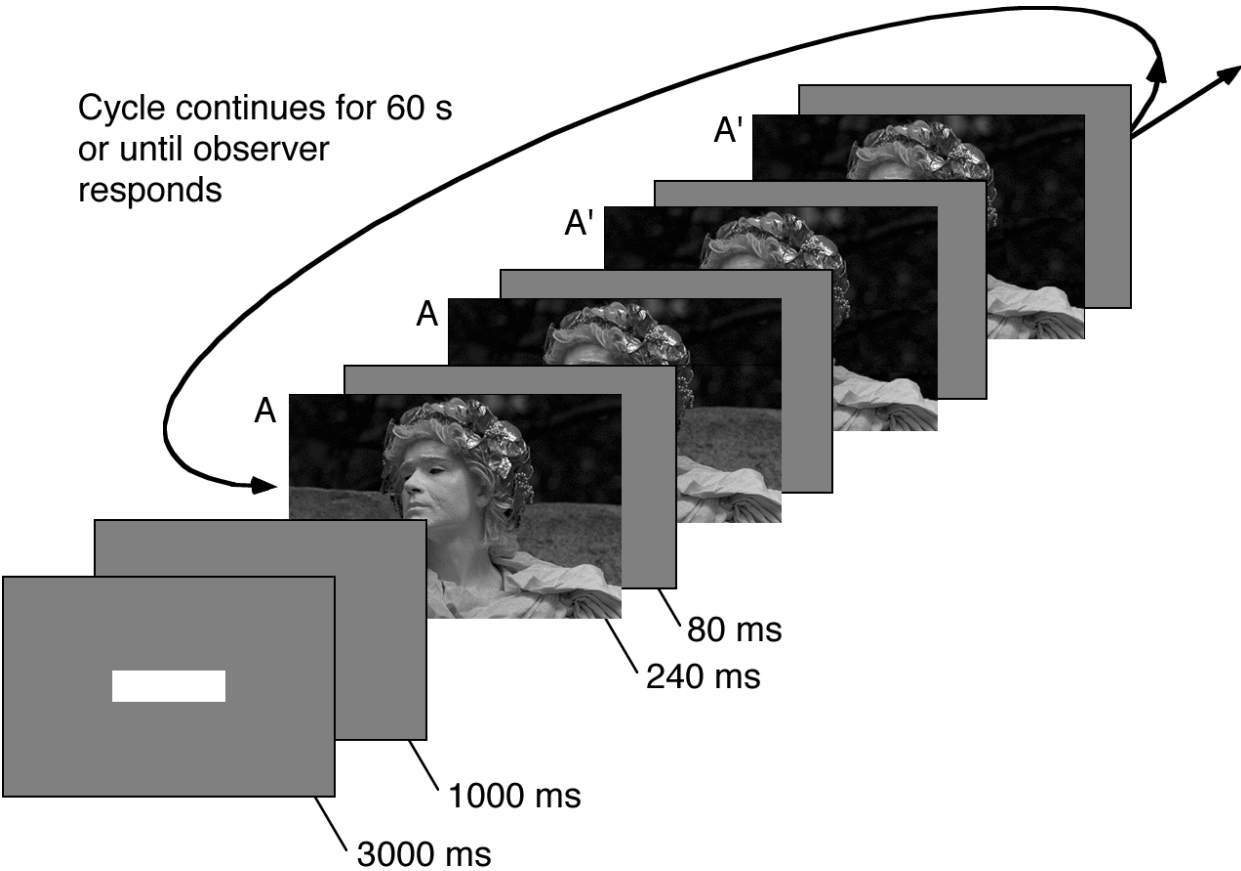
Figure 1

(a) Change in marginal interest (MI)

(b) Change in central interest (CI)

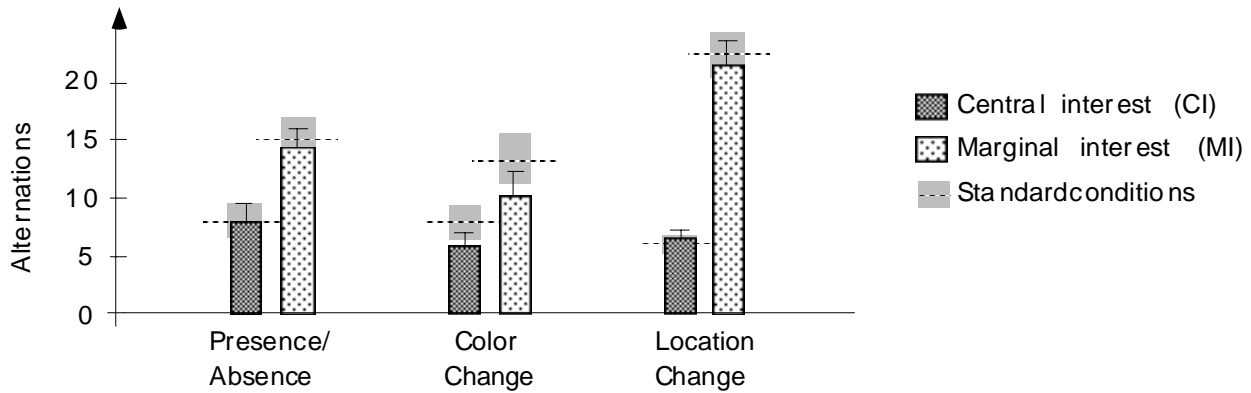**Figure   2**
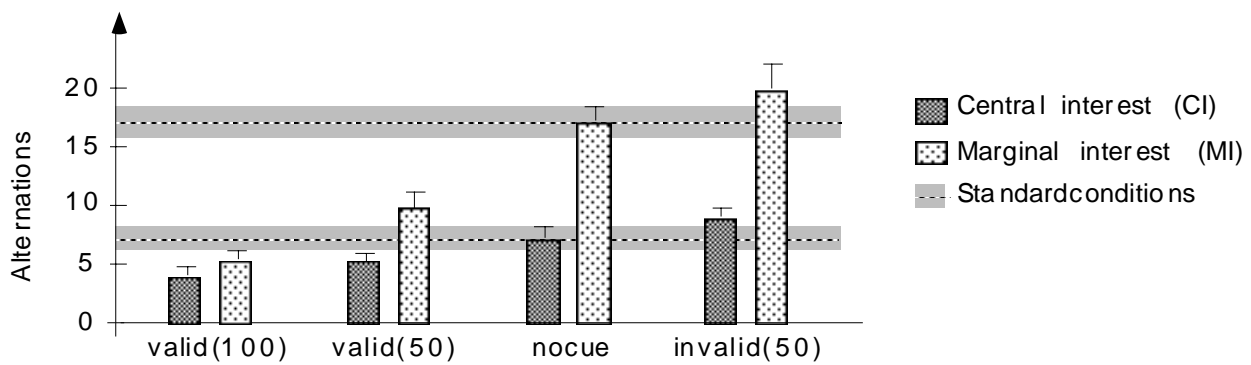
(a) - Standard flicker condition

(b) - Extended display duration

(c) - Effects of cue validity

Figure 3